

SYSTEM RESOURCE AVAILABILITY MANAGER

ABSTRACT OF THE DISCLOSURE

5 A hierarchical, distributed Availability Management (AM) process for
recovering from component failures in a data processing system. The hierarchy of AM
elements track a failure modality hierarchy of the data processing system components.
For example, the system hierarchy may include system cards, processors, and processes,
in which case the associated AM elements may be implemented at a card manager (CM)
10 level, a system manager (SM) level, and a process manager (PM) level. The AM
hierarchy is designed to achieve a failure granularity so that failures in the lower levels
of the hierarchy have less of an impact on the entire system. Each AM element is
responsible for receiving failure notifications from processing system components
associated with a next lower level of the hierarchy. Upon such indication, if the AM
15 element determines that the failed component may be restarted, if the failed component
may be restarted, the AM element then determines if it can be hot, warm, or cold
restarted and it does so without further notification or implication to system availability
of other components. Hot restart requires complete integrity of state information, warm
restart causes a recovery of last known good state information, and a cold restart results
20 in the re-initialization of state information. If, the component cannot be restarted, then
notification is provided to the next higher level of the hierarchy and the AM element
itself terminates. One of the AM processes may execute as an identity management
protocol. The identity protocol sets a temporary master state; waits a predetermined
amount of time; and then sets a final master state only if no other system card has
25 asserted a temporary master state. The waiting time period is selected to be greater than
the longest expected initialization process for peer components in the system.